arXiv:1607.02447v1 [q-bio.QM] 8 Jul 2016

A Weighted Exact Test for Mutually Exclusive Mutations in Cancer

Mark D.M. Leiserson, Matthew A. Reyna, Benjamin J. Raphael*
Department of Computer Science and
Center for Computational Molecular Biology
Brown University
Providence, 02912, USA

July 11, 2016

Abstract

Motivation: The somatic mutations in the pathways that drive cancer development tend to be mutually exclusive across tumors, providing a signal for distinguishing driver mutations from a larger number of random passenger mutations. This mutual exclusivity signal can be confounded by high and highly variable mutation rates across a cohort of samples. Current statistical tests for exclusivity that incorporate both per-gene and per-sample mutational frequencies are computationally expensive and have limited precision.

Results: We formulate a weighted exact test for assessing the significance of mutational exclusivity in an arbitrary number of mutational events. Our test conditions on the number of samples with a mutation as well as per-event, per-sample mutation probabilities. We provide a recursive formula to compute p-values for the weighted test exactly as well as a highly accurate and efficient saddlepoint approximation of the test. We use our test to approximate a commonly used permutation test for exclusivity that conditions on per-event, per-sample mutation frequencies. However, our test is more efficient and it recovers more significant results than the permutation test. We use our Weighted Exclusivity Test (WExT) software to analyze hundreds of colorectal and endometrial samples from The Cancer Genome Atlas, which are two cancer types that often have extremely high mutation rates. On both cancer types, the weighted test identifies sets of mutually exclusive mutations in cancer genes with fewer false positives than earlier approaches.

Availability: See `http://compbio.cs.brown.edu/projects/wext` for software.

Contact: braphael@cs.brown.edu

1 Introduction

A key challenge in cancer genomics is distinguishing the small number of somatic mutations that drive cancer from the vast majority of mutations that accumulate randomly. The ability to distinguish these *driver* mutations from the random *passenger* mutations may lead to better understanding of cancer biology and personalized therapies customized to a tumor's mutational profile. However, large scale cancer sequencing efforts such as The Cancer Genome Atlas (TCGA) [29, 31, 30] and the International Cancer Genome Consortium (ICGC) have shown that many driver mutations are rare across patient cohorts and thus distinguishing the driver mutations from the passengers by their frequency of occurrence is a difficult problem.

Driver mutations are hypothesized to group into a small number of pathways or hallmarks [10], and this hypothesis is a widely accepted explanation for the observed mutational heterogeneity of cancer [35]. Thus, researchers have developed methods to identify combinations of mutations using varying levels of prior knowledge, from pathway databases [37, 24, 27] to protein-protein interaction networks [33, 7, 26, 17].

*To whom correspondence should be addressed.

Because prior knowledge of pathways and interactions is often noisy or unavailable, *de novo* methods that do not use prior information are advantageous. The vast number of possible combinations of mutated genes makes complete de novo discovery of combinations computationally and statistically intractable. However, a number of methods [22, 34, 7, 16, 28, 8, 4, 18, 13, 14] use the observation that mutations within the same pathway are often mutually exclusive across tumors [32, 39]. These methods differ in how they score mutual exclusivity and in how they identify the best scoring set(s) of mutations.

The first type of score for mutual exclusivity is a combinatorial score, such as the scores employed in [22, 34, 16]. For example, in the Dendrix algorithm [34], the score for a set M of mutational events is the difference between the number of samples with a mutation in M (coverage) and the number of mutations in M occurring in more than one sample (coverage overlap). The advantage of a combinatorial score is that it is easy to compute, but it was observed by [18] and others that the score is often biased towards sets with frequently mutated genes.

The second type of score for mutual exclusivity is a statistical score [7, 19, 28, 8, 18, 13, 4, 14]. A particularly useful statistical score for exclusivity is based on the exact distribution that conditions on the observed number of mutated samples in each gene [18, 4]. For a pair of mutations, such a test is a one-sided Fisher's exact test for independence [19, 18, 4]. For more than two genes, [18] generalized the exact test to multi-dimensional contingency tables. They introduced the CoMEt algorithm that computes a generalization of Fisher's exact test for event sets of any size using either an exact tail enumeration algorithm or an approximation. They showed that conditioning on the number of mutations in each events reduces bias towards frequently mutated events compared to combinatorial scores.

Statistical scores that condition only on mutation frequencies do not account for the variation in mutation rate among tumors. It has been observed that the number of mutations in a tumor can vary over several orders of magnitude (e.g., see [35, 15, 25]). For example, colorectal tumors with microsatellite stability have a median of 66 nonsynonymous mutations, but colorectal tumors with microsatellite instability have a median of 777 mutations [35]. Another example is from [31], who classified a subset of TCGA endometrial cancers as ultramutated or hypermutated.

Another useful statistical test for mutual exclusivity conditions on *both* the number of mutated samples in each event *and* the number of mutated events in each sample [7, 13]. Since computing this distribution exactly is not computationally efficient, permutation tests are used. The permutation tests, which compare observed results to a number of samples ($\sim 10^4$) drawn from a null distribution, are more tractable than computing the p-value exactly on genome-scale data, but the significance of the score is directly limited by the number of permutations. MEMo [7] computes the significance of the *coverage* (number of mutated samples) of M using this permutational distribution for sets of any size k. MEMCover [13] computes the significance of the exclusivity of pairs to search for exclusivity within, between, and across cancer types. Both MEMo and MEMCover restrict their analysis to sets of genes that interact in a protein-interaction network. WeSME [14], which appeared while this paper was under review, computes the significance of exclusivity of pairs of genes with a less expensive approximation to the permutational distribution. To our knowledge, there is no method for quickly computing the significance of mutual exclusivity conditioned on both the observed number of mutations per event and number of mutations per sample.

1.1 Contributions

We introduce a weighted test for mutual exclusivity that conditions on the frequency of each mutational event in a set M and also incorporates the probability that each event is mutated in each sample. Our test was inspired by a model derived by [20], who computed the significance of the overlap between two sets of genes weights by gene length. We introduce the weighted exclusivity test to approximate the fixed gene and sample frequency permutation test quickly and accurately by estimating the mutation probabilities from the null distribution of the permutation test. We present a recursive formula for computing the p-value of this test exactly and derive a saddlepoint approximation for arbitrarily sized groups of genes. We show that the saddlepoint approximation is both a fast and accurate approximation of the permutational distribution. We also demonstrate that the saddlepoint approximation can be used to rapidly compute the CoMEt statistical test, which is a special case of the weighted test where the mutation probabilities for a given gene are the same in each sample.

We use our Weighted Exclusivity Test (WExT) software to identify sets of exclusive mutations in hundreds of colorectal and endometrial cancers. Cancer of these types often have extremely high mutation rates (e.g., see [35]), which makes them difficult to analyze when conditioning only on the number of mutated samples per event. However, our weighted statistical test allows us to effectively condition on the number of mutated genes per sample, and we identify exclusive patterns of mutations in these

Test	p-value	Conditioning	Algorithms
Row (R) exclusivity	Φ_{R}	Event frequencies	Tail enumeration (CoMEt), **saddlepoint approximation**
Row-column (RC) exclusivity	Φ_{RC}	Event and sample frequencies	Permutations
Weighted-row (WR) exclusivity	Φ_{WR}	Event frequencies and **per-event, per-sample weights (W)**	**Recursive formula**, **saddlepoint approximation**

Table 1: Three tests for mutual exclusivity, the values that are fixed in each test, and different algorithms for computing the p-values associated with the tests. **Boldface** entries indicate contributions by this manuscript.

cancers that were missed by earlier approaches. We find that the weighted test identifies more biologically interesting sets than CoMEt [18]. We expect that the weighted test for mutual exclusivity will prove useful for many cancer types where defects in DNA damage or environmental exposures, e.g., ultraviolet light, lead to very high mutation rates in some samples.

2 Methods

We introduce a new weighted test for mutual exclusivity that incorporates per-event, per-sample mutation probabilities, and we describe how to use particular instances of our test to approximate commonly used tests for mutual exclusivity, which we refer to as the *row exclusivity (R-exclusivity)* and *row-column exclusivity (RC-exclusivity)* tests.

First, in Section 2.2, we describe the R-exclusivity and RC-exclusivity tests. Next, in Section 2.3, we introduce our new weighted test for mutual exclusivity, which we call the *weighted-row exclusivity (WR-exclusivity) test*, that incorporates event and sample mutation frequencies without using permutations. In Section 2.4, we describe how to approximate the R-exclusivity and RC-exclusivity tests with the WR-exclusivity test. Then, in Section 2.5, we provide a recursive formula for computing the WR-exclusivity p-value exactly, and we derive a fast and accurate saddlepoint approximation for the WR-exclusivity p-value. Finally, in Sections 2.6, we describe how we search for exclusive sets, and in Section 2.7, we describe our WExT software.

We summarize the tests and our contributions in this paper in Table 1.

2.1 Notation

We observe the presence or absence of mutational events across a collection of samples. The presence of an event may reflect a variety of genomic (e.g., the canonical *BRAF* V600E mutation or deletions in *CDKN2A*), proteomic, and/or epigenomic alterations. In this work, we analyze single nucleotide variants and small insertion/deletions (indels) by gene. For the clarity of exposition, we will describe these events at the gene level, but our weighted test can accommodate a broader class of mutational events.

Let $\{g_i\}_{i=1}^m$ be a set of m genes and $\{s_j\}_{j=1}^n$ be a set of n samples. For each sample, we observe the presence of one or more mutations in each gene, and we record the presence or absence of mutations in a per-gene, per-sample binary mutation matrix $A \in \{0,1\}^{m \times n}$, where $A = [a_{ij}]$ with $a_{ij} = 1$ if gene g_i is mutated in sample s_j and $a_{ij} = 0$ otherwise.

Let $M \subseteq \{g_i\}_{i=1}^m$ be a set of k genes. The gene set M has *co-occurring* mutations in sample s_j if multiple genes are mutated in that sample, i.e., there exist distinct $g_i, g_\ell \in M$ such that $a_{ij} = 1$ and $a_{\ell j} = 1$. Alternatively, the gene set M has a *mutually exclusive* mutation in sample s_j if one and only one gene is mutated in that sample, i.e., there exists $g_\ell \in M$ such that $a_{\ell j} = 1$ and $a_{ij} = 0$ for $g_i \in M \setminus \{g_\ell\}$. Our goal is to identify sets of genes with statistically significant numbers of mutually exclusive mutations.

Let $r_i = \sum_{j=1}^n a_{ij}$ be the number of samples with mutations in g_i, let $c_j = \sum_{i=1}^m a_{ij}$ be the number of genes with mutations in s_j, let z_M be the number of samples with co-occurring mutations in M, and let t_M be the number of samples with mutually exclusive mutations in M.

For any mutation matrix B, let $B(M)$ be the submatrix of B with rows corresponding to the gene set M, and let $t_{B(M)}$ be the number of mutually exclusive mutations in $B(M)$. We will use $t_M = t_{A(M)}$.

2.2 Permutation tests for mutual exclusivity

We describe two different permutation tests for mutual exclusivity. First, the *row-exclusivity (R-exclusivity)* test finds the probability $\Phi_{\mathrm{R}}(M)$ of observing at least t_M mutually exclusive mutations in a gene set M given that each $g_i \in M$ is mutated in r_i samples. We describe this test as the row-exclusivity test because it conditions on the row sums of the mutation matrix.

Formally, we define Ω_{R} to be the set of mutation matrices with same row sums as A. Let $\mathcal{E}_{\mathrm{R}} = \left\{B \in \Omega_{\mathrm{R}} : t_{B(M)} \geq t_M\right\}$ be the set of mutation matrices with at least t_M mutually exclusive mutations in M. Then

$$\Phi_{\mathrm{R}}(M) = \frac{|\mathcal{E}_{\mathrm{R}}|}{|\Omega_{\mathrm{R}}|} \tag{1}$$

is the p-value of the R-exclusivity test.

Since the R-exclusivity test only conditions on the row sums of A, we can consider each row of A independently. This implies that to compute $\Phi_{\mathrm{R}}(M)$, we use only the rows corresponding to M. Thus, for $k = 2$, the p-value $\Phi_{\mathrm{R}}(M)$ is equal to the p-value from the one-sided Fisher's exact test, which computes the tail probability by summing the hypergeometric probability of 2×2 contingency tables with fixed margins. The hypergeometric probability of each contingency table is the proportion of matrices in Ω_{R} that give a contingency table with those margins. Note also that, when $k = 2$, the probability of observing t_M or more mutually exclusive mutations is equal to the probability of observing z_M or more co-occurring mutations. [18] generalized this test to $k > 2$ genes as part of the CoMEt algorithm.

The *row-column-exclusivity (RC-exclusivity)* test finds the probability $\Phi_{\mathrm{RC}}(M)$ of observing at least t_M mutually exclusive mutations in a gene set M given that each $g_i \in M$ is mutated in r_i samples *and* each s_j is mutated in c_j genes. We describe this test as the row-column-exclusivity test because it conditions on the row *and* column sums of the mutation matrix.

Formally, we define Ω_{RC} to be the set of mutation matrices with the same row and column sums as A. Let $\mathcal{E}_{\mathrm{RC}} = \left\{B \in \Omega_{\mathrm{RC}} : t_{B(M)} \geq t_M\right\}$ be the set of mutation matrices with at least t_M mutually exclusive mutations in M. Then

$$\Phi_{\mathrm{RC}}(M) = \frac{|\mathcal{E}_{\mathrm{RC}}|}{|\Omega_{\mathrm{RC}}|} \tag{2}$$

is the p-value of the RC-exclusivity test. Since Ω_{RC} depends on the row and column sums of A, we cannot consider the rows of A, or even $A(M)$, independently.

The RC-exclusivity test is related to the co-occurrence and mutual exclusivity tests used in [7] and [13, 14] with a few key differences. First, [7] use coverage (i.e., $t_M + z_M$) instead of exclusivity as the test statistic, while [13, 14] limit to pairs of genes. Second, both [7] and [13] use permutation tests that sample matrices from Ω_{RC}, so their p-values are limited by the number of draws (e.g., both use 10^4 permutations).

2.3 Weighted exact test for mutual exclusivity

We introduce a new weighted test for mutual exclusivity. The *weighted-row-exclusivity (WR-exclusivity)* test finds the probability $\Phi_{\mathrm{WR}}(M)$ of observing at least t_M mutually exclusive mutations in a gene set M given that $g_i \in M$ is mutated in r_i samples *and* a per-gene, per-sample mutation probability matrix W that prescribes weights with the presence or absence of individual mutations. We describe this test as the weighted-row-exclusivity test because it conditions on the row sums of the mutation matrix and a probability weight matrix.

For our model, we assume that $\{X_{ij}\}_{j=1}^{n}$ is a set of mutually independent Bernoulli trials for each gene g_i with success probabilities $W = [w_{ij}]$, i.e.,

$$\Pr(X_{ij} = \ell) = \begin{cases} w_{ij}, & \text{if } \ell = 1, \\ 1 - w_{ij}, & \text{if } \ell = 0, \end{cases} \tag{3}$$

where w_{ij} is the probability that gene g_i is mutated in sample s_j. Let $T_{M,j}$ be a random variable with $T_{M,j} = 1$ if s_j has a mutually exclusive mutation in a gene set M and $T_{M,j} = 0$ otherwise. Therefore, $Y_i = \sum_{j=1}^{n} X_{ij}$ is a Poisson binomial distributed variable for the number of mutations in g_i and $T_M = \sum_{j=1}^{n} T_{M,j}$ is a test statistic for mutual exclusivity indicating the number of mutually exclusive mutations in M. We want to find the tail probability (commonly referred to as the p-value) of observing at least t_M mutually exclusive mutations in M given that g_i is mutated in r_i samples. The WR-exclusivity p-value $\Phi_{\mathrm{WR}}(M)$ is the probability of observing at least t_M mutations in a gene set M under this model with

$$\Phi_{\mathrm{WR}}(M) = \Pr(T_M \geq t_M \mid Y_M = \mathbf{r}_M) \tag{4}$$

where $Y_M = [Y_i]_{i \in M}$ and $\mathbf{r}_M = [r_i]_{i \in M}$.

Note that, for any gene g_i, the assumption that $Y_i = r_i$ implies that

$$\sum_{j=1}^{n} w_{ij} = \sum_{j=1}^{n} E[X_{ij}] = E\left[\sum_{j=1}^{n} X_{ij}\right] = E[Y_i] = r_i \tag{5}$$

by the definitions of $\{X_{ij}\}_{j=1}^{n}$ and Y_i.

2.4 Approximating the permutation tests with the weighted exclusivity test

Each of the sets Ω_{R} and Ω_{RC} underlying the R-exclusivity and RC-exclusivity tests, respectively, determines a per-gene, per-sample weight matrix $W = [w_{ij}]$ by considering the probability w_{ij} of observing a mutation in gene g_i in sample s_j, i.e.,

$$W = \frac{1}{|\Omega|} \sum_{B \in \Omega} B \tag{6}$$

where $\Omega \in \{\Omega_{\mathrm{R}}, \Omega_{\mathrm{RC}}\}$. Since both Ω_{R} and Ω_{RC} fix the number of mutated samples per gene, the weight matrix W in (6) with $\Omega \in \{\Omega_{\mathrm{R}}, \Omega_{\mathrm{RC}}\}$ satisfies (5). We define W_R to be the weight matrix with $\Omega = \Omega_{\mathrm{R}}$ and W_{RC} to be the weight matrix with $\Omega = \Omega_{\mathrm{RC}}$.

For the R-exclusivity test, each row of $B \in \Omega_{\mathrm{R}}$ can be considered separately, so (6) for the set Ω_{R} is given by $W_R = [w_{ij}]$, with $w_{ij} = \frac{r_i}{n}$.

However, for the RC-exclusivity test, each row $B \in \Omega_{\mathrm{RC}}$ cannot be considered separately, so, to the best of our knowledge, there is no closed-form expression for (6) for the set Ω_{RC}. Therefore, we generate an empirical weight matrix $W_{RC}^{N} = [w_{ij}]$ for Ω_{RC} by drawing N matrices Ω_{RC}^{N} uniformly at random from Ω_{RC} and computing (6) with Ω_{RC}^{N} instead of Ω_{RC}. We assume that there is a nonzero probability that any gene can be mutated in any sample, and thus set $w_{ij} = \frac{1}{2N}$ when no mutation in gene g_i is observed in sample s_j in Ω_{RC}^{N}.

Estimating W_{RC}^{N} in this way gives an accurate approximation of $\Phi_{\mathrm{RC}}(M)$ using relatively small values of N.

2.5 Computing the weighted exclusivity test

Our weighted test for mutual exclusivity requires computing the tail probability in (4), which can be computationally expensive. We compute the tail probability using two different strategies: a recursive formula and a saddlepoint approximation.

2.5.1 Recursive formula for the weighted exclusivity test

We present a recursive formula for computing the tail probability in (4) exactly for sets M of any size k. Assuming that $\{Y_i\}_{i=1}^{m}$ are mutually independent, we can write (4) as

$$\Phi_{\mathrm{WR}}(M) = \frac{\Pr(T_M \geq t_M, Y_M = \mathbf{r}_M)}{\prod_{i \in M} \Pr(Y_i = r_i)}. \tag{7}$$

Without loss of generality, let $M = \{1, \ldots, k\}$. We first find the joint probability in the numerator of (7) using a recursive formula, where $\Pr(T_M \geq t_M, Y_M = \mathbf{r}_M) = F(t_M, r_1, \ldots, r_k, n)$ is computed by the recurrence relation

$$\begin{aligned} &F(t, x_1, \ldots, x_k, j) = \\ &\sum_{\pi \in \{0,1\}^k} \prod_{i=1}^{k} q_{ij\pi_i} F(w_\pi(t), y_{\pi_1}(x_1), \ldots, y_{\pi_k}(x_k), j-1), \end{aligned} \tag{8}$$

where

$$q_{ij\ell} = \begin{cases} p_{ij} & \text{if } \ell = 1, \\ 0 & \text{otherwise,} \end{cases}$$

$$w_\pi(t) = \begin{cases} t-1 & \text{if } \sum_{i=1}^{k} \pi_i = 1, \\ t & \text{otherwise,} \end{cases}$$

and

$$y_\ell(x) = \begin{cases} x-1 & \text{if } \ell = 1, \\ x & \text{otherwise.} \end{cases}$$

The base cases for (8) are

$$F(t, x_1, \ldots, x_k, j) = \begin{cases} 1, & \text{if } t = x_1 = \cdots = x_k = j = 0, \\ 0, & \text{if } \min\{t, x_1, \ldots, x_k, j\} < 0, \\ & \quad t > \sum_{i=1}^{k} x_i, \text{ or} \\ & \quad \max_{i=1}^{k} x_i > n. \end{cases} \tag{9}$$

We then find the marginal probabilities in the denominator of (7) using dynamic programming, which is a standard method for computing the Poisson-Binomial probability mass function [12].

2.5.2 Saddlepoint approximation for the weighted exclusivity test

We derive a saddlepoint approximation [6] for computing the tail probability in (4). This approach is inspired by [20], who derive a saddlepoint approximation for an enrichment test for differentially expressed genes in Gene Ontology categories weighted by gene lengths. The saddlepoint approximation is specifically designed to provide a quick and accurate approximation of the tail probability. We present the key equation in (10) and provide a fuller derivation for $k = 3$ in the supplement. The saddlepoint approximation is given by

$$\Pr(T_M \geq t_M \mid Y_M = \mathbf{r}_M) \approx 1 - \Phi(\tilde{w}) - \phi(\tilde{w}) \left(\frac{1}{\tilde{w}} - \frac{1}{\tilde{u}} \right), \tag{10}$$

where Φ and ϕ are, in this setting, the cumulative distribution and density functions, respectively, of the standard normal distribution, and $\tilde{w}$ and $\tilde{u}$ are defined as follows.

Without loss of generality, let $M = \{1, \ldots, k\}$. First, for $\lambda \in \mathbb{R}^{k+1}$, let $\mathcal{M}_{Y_M, T_M}(\lambda) = E[e^{\sum_{i \in M} \lambda_i Y_i + \lambda_{k+1} T_M}]$ be the joint moment generating function of $\{Y_i\}_{i \in M}$ and T_M, and let $\mathcal{K}_{Y_M, T_M}(\lambda) = \log \mathcal{M}_{Y_M, T_M}(\lambda)$ be the corresponding joint cumulant generating function. Similarly, let $\mathcal{M}_{Y_i}(\lambda) = E[e^{\lambda_i Y_i}]$ be the moment generating function of Y_i, and let $\mathcal{K}_{Y_i}(\lambda) = \log \mathcal{M}_{Y_i}(\lambda)$ be the corresponding cumulant generating function.

Next, let $\mathcal{K}'_{Y_M, T_M}(\lambda)$ and $\mathcal{K}''_{Y_M, T_M}(\lambda)$ be the gradient vector and Hessian matrix, respectively, of $\mathcal{K}_{Y,T}(\lambda)$, and let $\mathcal{K}'_{Y_i}(\lambda)$ and $\mathcal{K}''_{Y_i}(\lambda)$ be the gradient vector and Hessian matrix, respectively, of $\mathcal{K}_{Y_i}(\lambda)$.

Finally, define $\tilde{w}$ by

$$\tilde{w} = \sqrt{2}\, \mathrm{sgn}\,(\tilde{y}_{k+1}) \sqrt{\sum_{i \in M} \mathcal{K}_{Y_i}(\hat{x}_i) - \mathcal{K}_{Y_M, T_M}(\tilde{y}) - \tilde{y}^T (\hat{x} - \tilde{x})} \tag{11}$$

and $\tilde{u}$ by

$$\tilde{u} = 2 \sinh\left(\frac{\tilde{y}_{k+1}}{2} \right) \sqrt{\frac{\left| \mathcal{K}''_{Y_M, T_M}(\tilde{y}) \right|}{\prod_{i \in M} \mathcal{K}''_{Y_i}(\hat{x}_i)}}, \tag{12}$$

where $\tilde{x} = \left(r_1, \ldots, r_k, t_M - \frac{1}{2}\right)$ and $\tilde{y} = (\tilde{y}_1, \ldots, \tilde{y}_{k+1})$ with $\tilde{y}$ the unique solution for $\mathcal{K}'_{Y_M, T_M}(\tilde{y}) = \tilde{x}$ and (10) undefined if $\tilde{y}_{k+1} = 0$, and $\hat{x} = (\hat{x}_1, \ldots, \hat{x}_k, 0)$ with $\hat{x}_i$ the unique solution for $\mathcal{K}'_{Y_i}(\hat{x}_i) = r_i$.

2.6 Searching for sets of mutually exclusive mutations

Our goal is to identify sets M of genes with significantly exclusive mutations, i.e., extremely small p-values $\Phi_{\mathrm{WR}}(M)$. There has been a considerable amount of work on methods for optimizing scores for mutually exclusive mutations, including Markov chain Monte Carlo methods [34, 18], integer linear programs [16, 40], greedy algorithms [4], and others. These methods have been shown to be able to search datasets of many hundreds of mutation events for mutually exclusive mutations. Many of these methods can be modified to use our weighted exclusivity test to identify the most significant sets.

Since our focus is a statistical test for exclusivity, we instead enumerate all sets M of k genes satisfying the following basic criteria and test them with the R-exclusivity, RC-exclusivity, and WR-exclusivity tests:

1. The number t_M of samples with mutually exclusive mutations must be larger than the number z_M of samples with co-occurring mutations, i.e., $t_M > z_M$.

2. Each gene $g_i \in M$ must have at least one exclusive mutation.

We use the Benjamini-Hochberg procedure [5] to control the false discovery rate (FDR). We examine the subset of genes in each dataset with a minimum mutation frequency so that we can enumerate and test all combinations of genes of a certain size in a reasonable amount of time.

2.7 Implementation

We implemented the recursive formula for the WR-exclusivity test in Python and C, and we implemented the saddlepoint approximation for the WR-exclusivity test in Python using the NumPy and SciPy numerical libraries. We implemented the RC-exclusivity test in Python, and we used a bipartite double edge swap algorithm (see [23, 9]) that has been shown empirically to sample uniformly from Ω_{RC}. Our code, along with commands and data for reproducing the results and figures in this paper, is available as the *Weighted Exclusivity Test (WExT)* software package at `http://compbio.cs.brown.edu/projects/wext`.

3 Results

We compare the results of the WR-exclusivity test to both the R-exclusivity and RC-exclusivity tests on real data. In general, we can choose any weights to compute WR-exclusivity, but, in this paper, we specifically consider weights to allow us to approximate the R-exclusivity and RC-exclusivity tests. We use WExT to discover mutually exclusive sets of mutations in thyroid, colorectal, and endometrial cancers, restricting our analysis to mutations at the gene level.

The rest of this section is organized as follows. In Section 3.1, we describe the data used in our experiments. In Section 3.2, we compare the tail enumeration and saddlepoint approximation algorithms for computing the R-exclusivity p-values $\Phi_{\mathrm{R}}(M)$, and we show that the saddlepoint approximation provides a fast and accurate approximation for $\Phi_{\mathrm{R}}(M)$. In Section 3.3, we compare the results of the recursive and saddlepoint approximation algorithms for computing the WR-exclusivity p-values $\Phi_{\mathrm{WR}}(M)$ with the results of the RC-exclusivity test, and we show that $\Phi_{\mathrm{WR}}(M)$ is an accurate approximation of $\Phi_{\mathrm{RC}}(M)$ using either the recursive or saddlepoint approximation algorithms. In Section 3.4, we show that $\Phi_{\mathrm{WR}}(M)$ provides an accurate approximation of $\Phi_{\mathrm{RC}}(M)$ even with coarser estimates of the weight matrix W. Finally, in Sections 3.5 and 3.6, we present the results of the WR-exclusivity test on thyroid, colorectal, and endometrial cancers.

3.1 Data

We analyzed non-synonymous single nucleotide variants (SNVs) and small insertions or deletions (indels) in 224 colorectal (COADREAD) [29], 402 papillary thyroid carcinoma (THCA) [30], and 248 uterine corpus endometrial carcinoma (UCEC) [31] samples from The Cancer Genome Atlas (TCGA). We analyzed the mutations in the COADREAD and UCEC samples from the TCGA Pan-Cancer project [36] by downloading the mutations in Mutation Annotation Format (MAF) from Synapse [1]. We downloaded the mutations in THCA from Firehose [2]. We restricted our analysis to non-synonymous mutations, ignoring mutations classified as "Silent", "Intron", "3'UTR", "5'UTR", "IGR", "lincRNA", and "RNA". We also downloaded lists of hypermutator samples for COADREAD and UCEC. We created a list of 35 hypermutator samples in COADREAD listed in [29] in their Supplementary Table 3, and 82 hypermutator samples in UCEC listed by [31] as samples labeled "POLE OR MSI" in their Supplementary Datafile S1.1. We restrict our analysis to genes mutated in at least 20, 5, and 30 samples in the COADREAD, THCA, and UCEC datasets, analyzing 76, 30, and 62 genes in each dataset, respectively.

In general, COADREAD samples have the most mutated genes (median: 78.5), with COADREAD hypermutators with mutations in at least an order of magnitude more genes than non-hypermutators (median for hypermutators: 797; median for non-hypermutators: 69). THCA samples have the fewest mutated genes per sample (median: 12), with no hypermutators, while UCEC has more mutated genes per sample (median: 57.5) with UCEC hypermutators mutated in approximately an order of magnitude more genes than non-hypermutators (median hypermutators: 355; median non-hypermutators: 43.5). See Supplementary Figure S1.

For each dataset, we estimated the weights W_{RC}^{N} using the permutation procedure described in Section 2.4 using $N = 10^3$ permutations. We show the weights for each dataset in Figure 1.

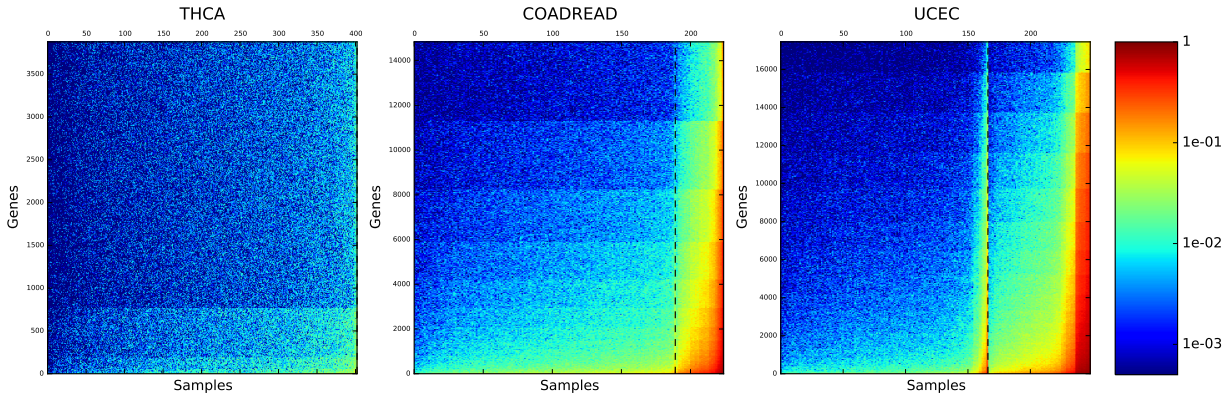

Figure 1: The weights W_N estimated by sampling $N = 10^3$ permuted matrices on the THCA, COADREAD, and UCEC datasets. Samples (x-axis) are sorted by the number of mutated genes in increasing order from left to right, with hypermutators (right) separated from non-hypermutators (left) with a dashed line in COADREAD and UCEC. Genes (y-axis) are sorted by the number of mutated samples in increasing order from top to bottom.

3.2 Comparison of methods for computing the R-exclusivity test on real data

First, we investigated the accuracy and speed of the saddlepoint approximation of the R-exclusivity p-value $\Phi_{\mathrm{R}}(M)$. We enumerated triples according to the procedure described in Section 2.6 in the THCA, COADREAD, and UCEC datasets, and computed $\Phi_{\mathrm{R}}(M)$ exactly using the CoMEt software from [18] as well as approximately using the saddlepoint approximation with W_R given in Section 2.4.

Supplementary Figure S2 shows a comparison of the p-values and runtimes given by the two methods, where the weights for the WR-exclusivity test are uniform across samples. On these datasets, the saddlepoint approximation is an extremely accurate approximation of the tail enumeration procedure ($\rho^2 = 0.995$). Additionally, while the median runtimes of the two algorithms are similar, the tail enumeration procedure is much slower for sets with co-occurring mutations while the saddlepoint approximation is largely unaffected. We expect the discrepancy between runtimes to grow for gene sets of larger sizes.

3.3 Comparison of methods for computing the WR-exclusivity test on real data

Next, we compared the results of methods for computing WR-exclusivity test with weights W_{RC}^N the RC-exclusivity test on pairs of genes from the THCA, COADREAD, and UCEC datasets. We chose pairs instead of triples because of the prohibitive cost of computing the recursive formula for $\Phi_{\mathrm{WR}}(M)$. We used $N = 10^4$ permutations to compute the $\Phi_{\mathrm{RC}}(M)$, and also included the tail enumeration procedure for $\Phi_{\mathrm{R}}(M)$ as a control.

Table 2 shows that the results of WR-exclusivity test – computed either with the recursive formula or the saddlepoint approximation – are strongly correlated with the RC-exclusivity test (Figure 2(b)). The results of the R-exclusivity test are more weakly correlated with the RC-exclusivity test (Figure 2(a)), showing that conditioning on the number of mutations in each sample changes the distribution of mutually exclusive mutations. This discrepancy remains when we restrict to gene sets M with $\Phi_{\mathrm{RC}}(M) \geq 10^{-4}$, i.e., sets of genes for which the empirical permutational distribution finds at least one mutually exclusive mutation in M.

The WR-exclusivity p-values computed exactly and with the saddlepoint approximation are highly correlated (Figure 2(c)), with a Pearson's correlation coefficient of 0.996 for all p. For smaller p-values with $\Phi_{\mathrm{WR}}(M) < 10^{-4}$ from either the recursive formula or the saddlepoint approximation, the correlation increases to 0.9999.

The runtime to compute Φ_{WR} using the recursive formula varies widely because pairs with co-occurring mutations require more computation, but the runtime of the saddlepoint approximation is more consistent. As a result, testing all pairs with the recursive formula requires approximately 2 hours,

Pairs	$\Phi_{\mathbf{R}}$ (CoMEt)	$\Phi_{\mathbf{WR}}$ (recursive)	$\Phi_{\mathbf{WR}}$ (saddlepoint)
All	0.71291	0.99816	0.99481
$\Phi_{\mathrm{RC}}(M) \geq 10^{-4}$	0.65376	0.99811	0.99404

Table 2: Pearson's correlation coefficient ρ^2 of p-values of pairs of genes from the THCA, COADREAD, and UCEC datasets using the tail enumeration R-exclusivity p-values $\Phi_{\mathrm{R}}(M)$, RC-exclusivity p-values $\Phi_{\mathrm{RC}}(M)$, and the recursive formula and saddlepoint approximations of the WR-exclusivity p-values $\Phi_{\mathrm{WR}}(M)$ using weights W_R. The correlations were computed for two sets of pairs of genes: p-values for 5,014 pairs (all) and 4,926 pairs ($\Phi_{\mathrm{RC}}(M) \geq 10^{-4}$).

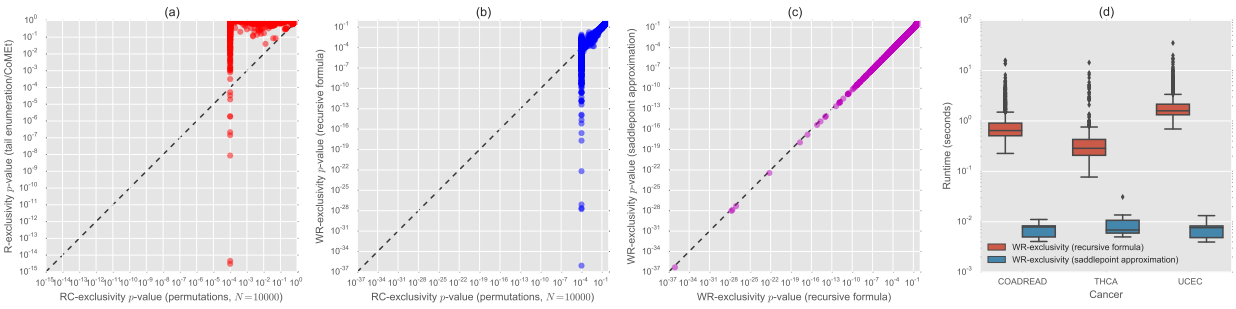

Figure 2: Comparison of p-values and runtimes of different tests on THCA, COADREAD, and UCEC pairs. (a-b) Scatter plots comparing the RC-exclusivity test with $N = 10^4$ permutations against (a) the R-exclusivity test and (b) the WR-exclusivity test (recursive) with weights W_{RC}^N. (c) The WR-exclusivity (recursive) p-values versus the WR-exclusivity (saddlepoint) p-values with weights W_{RC}^N. (d) Boxplots of the runtimes for computing the weighted test with the recusive formula (red) and with the saddlepoint approximation (blue) for each pair of genes in the datasets.

but testing the same pairs with the saddlepoint approximation requires approximately 30 seconds. Note that the runtime does not include generating the weights W_{RC}^N, which requires several minutes.

3.4 Approximating the RC-exclusivity test with the WR-exclusivity test

We compared the saddlepoint approximation of the WR-exclusivity test to the RC-exclusivity test using gene triples from the COADREAD dataset, again using the R-exclusivity test as a control. We computed $\Phi_{\mathrm{RC}}(M)$ with $N = 10^6$ permutations. We computed the saddlepoint approximation for $\Phi_{\mathrm{WR}}(M)$ using W_{RC}^N with $N = 10^3$ draws from Ω_{RC}, which is three orders of magnitude fewer than the number of permutations than we used to compute $\Phi_{\mathrm{RC}}(M)$. The p-values $\Phi_{\mathrm{R}}(M)$ and $\Phi_{\mathrm{RC}}(M)$ are weakly correlated in the tail ($\rho^2 = 0.67$ for $\Phi_{\mathrm{RC}}(M) < 0.001$; see Figure 3a). In contrast, the $\Phi_{\mathrm{WR}}(M)$ (saddlepoint) p-values provide an accurate approximation of the $\Phi_{\mathrm{RC}}(M)$ p-values. The RC-exclusivity and WR-exclusivity p-values are highly correlated in the tail ($\rho^2 = 0.948$ for $\Phi_{\mathrm{RC}}(M) < 0.001$; see Figure 3b). Moreover, $\Phi_{\mathrm{WR}}(M)$ is an accurate estimate of $\Phi_{\mathrm{RC}}(M)$ to within one or more digits for most triples and within an order of magnitude for all triples. Furthermore, despite the much smaller number of permutations used to generate W_{RC}^N, $\Phi_{\mathrm{WR}}(M)$ provides smaller p-values than $\Phi_{\mathrm{RC}}(M)$ with a greater number of significant predictions, and is much faster than the permutation test.

3.5 Mutually exclusive mutations in thyroid carcinomas

We computed the WR-exclusivity p-values for all triples of genes that were each mutated in at least 5 of the 402 thyroid carcinomas in the THCA dataset. The WR-exclusivity test identifies 48 triples with significantly exclusive mutations (FDR < 0.001), while the R-exclusivity test identifies 38 triples (FDR < 0.001).

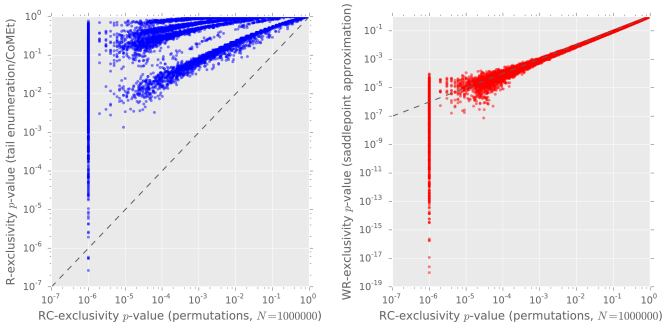

Figure 3: Comparison of the R-exclusivity test (left, y-axis) and the saddlepoint approximation of the WR-exclusivity test with weights W_{RC}^{N} estimated from $N = 10^3$ permutations (right, y-axis) to the R-exclusivity test (x-axis) on triples from the COADREAD dataset. The saddlepoint approximation uses weights W_{RC}^{N} estimated from $N = 10^3$ permutations, while the RC-exclusivity p-values were computed with $N = 10^6$ permutations.

The top 25 ranked triples by both tests are identical, which is not surprising since THCA samples have low mutation rates compared to most cancer types (see [35] and Section 3.1). In addition, the p-values for the top ranked triples are all within a few orders of magnitude, demonstrating that the two tests are very similar on this dataset.

Supplementary Table S1 shows the top triples, which include many known thyroid cancer genes. The top five triples include seven genes, five of which are well-known cancer genes with known roles in thyroid cancer [30]: *BRAF*, *HRAS*, *NRAS*, *EIF1AX*, and *ATM*. The other two genes are *BDP1* and *TG*, both of which may play a role in cancer. [38] describe a role for *BDP1* in AKT signaling, which was also noted in TCGA thyroid publication [30]), although *BDP1* is greater than 11,000 amino acids in length, so it may also accumulate many passenger mutations. *TG* is the thyroglobulin gene, and is used as a tumor marker in papillary thyroid carcinoma [3], which is the same subtype of thyroid cancer analyzed in TCGA.

3.6 Mutually exclusive mutations in colorectal cancers and endometrial carcinomas

We expect that the difference between the R-exclusivity and WR-exclusivity tests would be more pronounced on cancer types with higher and highly variable mutation rates. Thus, we computed p-values on triples of genes from colorectal cancers (COADREAD) and endometrial carcinomas (UCEC). We find that the WR-exclusivity test predicts more biologically interesting triples than the R-exclusivity test. The WR-exclusivity test identifies 5,290 and 6,835 triples (many of which overlap) with significantly mutually exclusive mutations (FDR < 0.001) in the 224 COADREAD and 248 UCEC samples, respectively. In contrast, the R-exclusivity test computes 4 and 130 triples, respectively, with significantly mutually exclusive mutations (FDR < 0.001).

Compared to the R-exclusivity test results, the highest ranked triples by the WR-exclusivity test include fewer long genes that tend to accumulate random, passenger mutations – especially in samples with high mutation rates (Tables 3 and 4).

On COADREAD, the WR-exclusivity test identifies ten different genes in the five most significant triples (Table 3). Nine of these genes are well-known cancer genes – *BRAF*, *KRAS*, *NRAS*, *ACV2RA*, *PIK3CA*, *TP53*, *ATM*, *TGFBR2*, and *ARID1A* – while the tenth gene (*ABCA12*) is known to have an association with colorectal cancers [11]. The R-exclusivity test results are similar – two of the top five triples identified by the WR-exclusivity test are in the top five triples identified by the R-exclusivity test – but the R-exclusivity test does not identify *ARID1A*, *TGFBR2*, *KRAS*, or *NRAS*. Further, the three additional genes identified by the R-exclusivity – *APC*, *FAT2*, and *WDFY3* – are all in the top 600 longest genes in the human genome (at least 9,560 amino acids). While mutations in *APC* are well-known to play a role in colorectal cancers, there is currently little evidence for the roles of *FAT2* or *WDFY3* in cancer, and it is likely that these long genes have accumulated many passenger mutations, particularly in hypermutated samples. Also of note is the fact that the number of hypermutator samples that contain mutations in the top triples from the WR-exclusivity test are not appreciably different from

$\Phi_{\mathbf{R}}$ rank	$\Phi_{\mathbf{WR}}$ rank	Triple M	$\Phi_{\mathbf{R}}(M)$	$\Phi_{\mathbf{WR}}(M)$	Hypermutator mutations
1	2	ACVR2A, PIK3CA, TP53	$2.65 \cdot 10^{-7}$	$2.54 \cdot 10^{-18}$	31
2	32	**APC**, BRAF, PRDM2	$5.44 \cdot 10^{-7}$	$2.30 \cdot 10^{-13}$	33
2	33	**APC**, BRAF, **WDFY3**	$5.44 \cdot 10^{-7}$	$2.43 \cdot 10^{-13}$	32
4	3	**ATM**, PIK3CA, TP53	$5.87 \cdot 10^{-7}$	$1.15 \cdot 10^{-17}$	24
5	81	**APC**, BRAF, **FAT2**	$1.93 \cdot 10^{-6}$	$6.48 \cdot 10^{-12}$	35
6	1	BRAF, KRAS, NRAS	$2.50 \cdot 10^{-6}$	$9.95 \cdot 10^{-19}$	26
1	2	ACVR2A, PIK3CA, TP53	$2.65 \cdot 10^{-7}$	$2.54 \cdot 10^{-18}$	31
4	3	**ATM**, PIK3CA, TP53	$5.87 \cdot 10^{-7}$	$1.15 \cdot 10^{-17}$	24
10	4	ARID1A, TGFBR2, TP53	$5.89 \cdot 10^{-6}$	$1.76 \cdot 10^{-16}$	29
9	5	ABCA12, TGFBR2, TP53	$4.29 \cdot 10^{-6}$	$1.83 \cdot 10^{-16}$	28

Table 3: Five most significant triples identified by the R-exclusivity (upper 5) and WR-exclusivity (lower 5) tests on the COADREAD dataset. Genes in bold are among 600 longest genes (at least 9,560 amino acids).

$\Phi_{\mathbf{R}}$ rank	$\Phi_{\mathbf{WR}}$ rank	Triple M	$\Phi_{\mathbf{R}}(M)$	$\Phi_{\mathbf{WR}}(M)$	Hypermutator mutations
1	20	**CACNA1E**, PTEN, TP53	$3.11 \cdot 10^{-12}$	$5.71 \cdot 10^{-30}$	77
2	21	**LAMA2**, PTEN, TP53	$4.13 \cdot 10^{-12}$	$8.05 \cdot 10^{-30}$	77
3	29	PTEN, **RYR2**, TP53	$4.60 \cdot 10^{-12}$	$4.85 \cdot 10^{-29}$	78
4	28	**NBEA**, PTEN, TP53	$8.40 \cdot 10^{-12}$	$3.32 \cdot 10^{-29}$	76
5	39	**FAT4**, PTEN, TP53	$1.23 \cdot 10^{-11}$	$3.0 \cdot 10^{-28}$	75
22	1	CTNNB1, RPL22, TP53	$2.11 \cdot 10^{-10}$	$1.20 \cdot 10^{-41}$	47
44	2	CTNNB1, KRAS, TP53	$3.05 \cdot 10^{-9}$	$1.10 \cdot 10^{-37}$	48
55	3	CTNNB1, MLL4, TP53	$4.26 \cdot 10^{-8}$	$5.78 \cdot 10^{-36}$	42
57	4	CTCF, CTNNB1, TP53	$4.84 \cdot 10^{-8}$	$3.29 \cdot 10^{-35}$	43
60	5	CTNNB1, **RYR1**, TP53	$1.14 \cdot 10^{-7}$	$9.51 \cdot 10^{-35}$	40

Table 4: Five most significant triples identified by the R-exclusivity (top 5) and WR-exclusivity (bottom 5) tests on the UCEC dataset. Notation as in Table 3.

the number of hypermutator samples that contain mutations in the top triples from the R-exclusivity test. This demonstrates that the WR-exclusivity is not systematically excluding hypermutator samples from consideration, but rather weighting the contribution of these samples appropriately in evaluating the significance of mutual exclusivity.

On UCEC, the differences between the R-exclusivity and WR-exclusivity tests are even more pronounced. The WR-exclusivity test identifies seven genes in the top five most significant triples (Table 4). These include six genes with known roles in cancer – *CTTNB1*, *TP53*, *RPL22*, *KRAS*, *CTCF*, and *MLL4* – with only one gene – *RYR1* – with likely spurious mutations. In contrast, the top five triples ranked by the R-exclusivity test include *PTEN*, and *TP53* – two well-known cancer genes – but also five genes with no known role in cancer that are all longer than 11,000 amino acids: *CACNA1E*, *LAMA2*, *RYR2*, *NBEA*, and *FAT4*. Further, none of the top five triples identified by the WR-exclusivity test are in the top twenty R-exclusivity triples. Finally, the R-exclusivity triples include many more mutations in hypermutator samples (ranging from mutations in 75 to 78 of the 81 hypermutators, versus 40 to 47 for the WR-exclusivity triples). This further demonstrates how the results of the R-exclusivity test are skewed by hypermutator samples, while the WR-exclusivity test incorporates the contribution of these samples appropriately in evaluating the significance of mutual exclusivity.

4 Discussion

We introduce a weighted exact test for the mutual exclusivity of mutations in cancer. We use this test to approximate the permutational test for exclusivity where the number of mutations in each event and each sample are fixed. To do so, we estimate per-event, per-sample mutation probabilities directly from the permutational distribution. We derive a recursive formula and a saddlepoint approximation of the p-value of the weighted test for event sets of any size, and we demonstrate the accuracy and efficiency of the saddlepoint approximation on genome-scale mutation datasets. Together, these contributions allow us to overcome the significant computational challenge of finding highly significant sets of mutually exclusive

mutations conditioned on both the observed number of mutations per event and per sample.

We then demonstrate the weighted test on three datasets with hundreds of samples from TCGA, including colorectal and endometrial cancers that have high variability in the number of mutations per sample. The weighted test identifies sets of mutually exclusive mutations including known cancer genes in each dataset, and its results include many fewer long genes and mutations in hypermutator samples than do the results of the generalization of Fisher's exact test from CoMEt [18].

There are several avenues for improving analyses with the weighted test. First, while we restricted our study to non-synonymous SNVs and indels, one should also test mutual exclusivity between other types of aberrations, such as copy number aberrations and gene fusions. We searched for mutually exclusive mutations by enumerating sets containing the most mutated genes, but the weighted test could easily be used in existing algorithms for optimizing mutual exclusivity scores (e.g., the MCMC from [34], the greedy approach from [4]) or to search for multiple sets simultaneously (e.g., from [18]). We estimated the per-event, per-sample mutation probability weights directly from the permutational distribution, but we also anticipate alternative methods for setting the weights that incorporate different event or sample attributes, such as gene length, to further reduce the number of false positives.

The weighted test may be of broader interest beyond searching for mutually exclusive mutations, both in other areas of computational biology and other disciplines. For example, statistical tests of "presence-absence" matrices with fixed row and column sums are a common tool in ecology for looking at species-associations, but can be computationally prohibitive (e.g., [21]). The weighted exact test presented here may offer a fast, alternative approach for computing the significance of associations with high accuracy.

5 Acknowledgments

The authors would like to acknowledge [20] for inspiring this work, and Uri Keich for generously helping us run their code and providing comments on our manuscript. This work is supported by US National Institutes of Health (NIH) grants R01HG005690, R01HG007069 and R01CA180776 to B.J.R. M.D.M.L. is supported by NSF fellowship GRFP DGE 0228243. B.J.R. is supported by a Career Award at the Scientific Interface from the Burroughs Wellcome Fund, an Alfred P. Sloan Research Fellowship and an NSF CAREER Award (CCF-1053753).

References

[1] `https://doi.org/10.7303/syn1710680.4`.

[2] `http://gdac.broadinstitute.org/runs/stddata__2016_01_28/data/THCA/20160128/gdac.broadinstitute.org_THCA.Mutation_Packager_Calls.Level_3.2016012800.0.0.tar.gz`.

[3] `http://www.mayomedicallaboratories.com/test-catalog/Clinical+and+Interpretive/62800`.

[4] Özgün Babur, Mithat Gönen, Bülent Arman Aksoy, Nikolaus Schultz, Giovanni Ciriello, et al. Systematic identification of cancer driving signaling pathways based on mutual exclusivity of genomic alterations. *Genome biology*, 16:45, 2015.

[5] Y Benjamini and Y Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B (Methodological)*, 57:289–300, 1995.

[6] Ronald W Butler. *Saddlepoint approximations with applications*, volume 22. Cambridge University Press, 2007.

[7] Giovanni Ciriello, Ethan Cerami, Chris Sander, and Nikolaus Schultz. Mutual exclusivity analysis identifies oncogenic network modules. *Genome research*, 22(2):398–406, 2012.

[8] Simona Constantinescu, Ewa Szczurek, Pejman Mohammadi, Jörg Rahnenführer, and Niko Beerenwinkel. TiMEx: a waiting time model for mutually exclusive cancer alterations. *Bioinformatics (Oxford, England)*, 2015.

[9] Andrea Gobbi, Francesco Iorio, Kevin J Dawson, David C Wedge, David Tamborero, et al. Fast randomization of large genomic datasets while preserving alteration counts. *Bioinformatics (Oxford, England)*, 30(17):i617–23, 2014.

[10] Douglas Hanahan and Robert A Weinberg. Hallmarks of cancer: the next generation. *Cell*, 144(5):646–674, 2011.

[11] I Hlavata, B Mohelnikova-Duchonova, R Vaclavikova, V Liska, P Pitule, P Novak, et al. The role of ABC transporters in progression and clinical outcome of colorectal cancer. *Mutagenesis*, 27(2):187–196, 2012.

[12] Yili Hong. On computing the distribution function for the Poisson binomial distribution. *Computational Statistics & Data Analysis*, pages 41–51, 2013.

[13] Yoo-Ah Kim, Dong-Yeon Cho, Phuong Dao, and Teresa M Przytycka. MEMCover: integrated analysis of mutual exclusivity and functional network reveals dysregulated pathways across multiple cancer types. *Bioinformatics (Oxford, England)*, 31(12):i284–92, 2015.

[14] Yoo-Ah Kim, Sanna Madan, and Teresa M Przytycka. Wesme: Uncovering mutual exclusivity of cancer drivers and beyond. *Bioinformatics (Oxford, England)*, page To appear, 2016.

[15] Michael S Lawrence, Petar Stojanov, Paz Polak, Gregory V Kryukov, Kristian Cibulskis, et al. Mutational heterogeneity in cancer and the search for new cancer-associated genes. *Nature*, 499(7457):214–218, 2013.

[16] Mark D M Leiserson, Dima Blokh, Roded Sharan, and Benjamin J Raphael. Simultaneous identification of multiple driver pathways in cancer. *PLoS Computational Biology*, 9(5):e1003054, 2013.

[17] Mark D M Leiserson, Fabio Vandin, Hsin-Ta Wu, Jason R Dobson, Jonathan V Eldridge, et al. Pan-cancer network analysis identifies combinations of rare somatic mutations across pathways and protein complexes. *Nature genetics*, 47(2):106–114, 2015.

[18] Mark D M Leiserson, Hsin-Ta Wu, Fabio Vandin, and Benjamin J Raphael. CoMEt: a statistical approach to identify combinations of mutually exclusive alterations in cancer. *Genome biology*, 16(1):160, 2015.

[19] Timothy J Ley, Christopher Miller, Li Ding, Benjamin J Raphael, Andrew J Mungall, et al. Genomic and Epigenomic Landscapes of Adult De Novo Acute Myeloid Leukemia. *NEJM*, 368(22):2059–2074, 2013.

[20] David Manescu and Uri Keich. A Symmetric Length-Aware Enrichment Test. *RECOMB*, pages 224–242, 2015.

[21] I Miklós and J Podani. Randomization of presence-absence matrices: comments and new algorithms. *Ecology*, 85(1):86–92, 2004.

[22] Christopher A Miller, Stephen H Settle, Erik P Sulman, Kenneth D Aldape, and Aleksandar Milosavljevic. Discovering functional modules by identifying recurrent and mutually exclusive mutational patterns in tumors. *BMC medical genomics*, 4:34, 2011.

[23] Ron Milo, Nadav Kashtan, Shalev Itzkovitz, Mark EJ Newman, and Uri Alon. On the uniform generation of random graphs with prescribed degree sequences. *arXiv preprint cond-mat/0312028*, 2003.

[24] Vamsi K Mootha, Cecilia M Lindgren, Karl-Fredrik Eriksson, Aravind Subramanian, Smita Sihag, et al. PGC-1alpha-responsive genes involved in oxidative phosphorylation are coordinately downregulated in human diabetes. *Nature genetics*, 34(3):267–273, 2003.

[25] Steven A Roberts and Dmitry A Gordenin. Hypermutation in human cancer genomes: footprints and mechanisms. *Nature Reviews Cancer*, 14(12):786–800, 2014.

[26] Matthew Ruffalo, Mehmet Koyutürk, and Roded Sharan. Network-Based Integration of Disparate Omic Data To Identify "Silent Players" in Cancer. *PLoS computational biology*, 11(12):e1004595, 2015.

[27] Aravind Subramanian, Pablo Tamayo, Vamsi K Mootha, Sayan Mukherjee, Benjamin L Ebert, et al. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences of the United States of America*, 102(43):15545–15550, 2005.

[28] Ewa Szczurek and Niko Beerenwinkel. Modeling mutual exclusivity of cancer mutations. *PLoS computational biology*, 10(3):e1003503, 2014.

[29] The Cancer Genome Atlas Research Network. Comprehensive molecular characterization of human colon and rectal cancer. *Nature*, 487(7407):330–337, 2012.

[30] The Cancer Genome Atlas Research Network. Integrated genomic characterization of papillary thyroid carcinoma. *Cell*, 159(3):676–690, 2014.

[31] The Cancer Genome Atlas Research Network, Cyriac Kandoth, Nikolaus Schultz, Andrew D Cherniack, Rehan Akbani, et al. Integrated genomic characterization of endometrial carcinoma. *Nature*, 497(7447):67–73, 2013.

[32] Roman K Thomas, Alissa C Baker, Ralph M Debiasi, Wendy Winckler, Thomas Laframboise, et al. High-throughput oncogene mutation profiling in human cancer. *Nature genetics*, 39(3):347–351, 2007.

[33] Fabio Vandin, Eli Upfal, and Benjamin J Raphael. Algorithms for detecting significantly mutated pathways in cancer. *Journal of computational biology : a journal of computational molecular cell biology*, 18(3):507–522, 2011.

[34] Fabio Vandin, Eli Upfal, and Benjamin J Raphael. De novo discovery of mutated driver pathways in cancer. *Genome research*, 22(2):375–385, 2012.

[35] Bert Vogelstein, Nickolas Papadopoulos, Victor E Velculescu, Shibin Zhou, Luis A Diaz, et al. Cancer genome landscapes. *Science*, 339(6127):1546–1558, 2013.

[36] John N Weinstein, Eric A Collisson, Gordon B Mills, Kenna R Mills Shaw, Brad A Ozenberger, et al. The Cancer Genome Atlas Pan-Cancer analysis project. *Nature Genetics*, 45(10):1113–1120, 2013.

[37] Michael C Wendl, John W Wallis, Ling Lin, Cyriac Kandoth, Elaine R Mardis, et al. PathScan: a tool for discerning mutational significance in groups of putative cancer genes. *Bioinformatics (Oxford, England)*, 27(12):1595–1602, 2011.

[38] Annette Woiwode, Sandra A S Johnson, Shuping Zhong, Cheng Zhang, Robert G Roeder, et al. PTEN represses RNA polymerase III-dependent transcription by targeting the TFIIIB complex. *Molecular and cellular biology*, 28(12):4204–4214, 2008.

[39] Chen-Hsiang Yeang, Frank McCormick, and Arnold Levine. Combinatorial patterns of somatic gene mutations in cancer. *FASEB journal : official publication of the Federation of American Societies for Experimental Biology*, 22(8):2605–2622, 2008.

[40] Junhua Zhang, Ling-Yun Wu, Xiang-Sun Zhang, and Shihua Zhang. Discovery of co-occurring driver pathways in cancer. *BMC bioinformatics*, 15:271, 2014.

Supplement for A Weighted Exact Test for Mutually Exclusive Mutations in Cancer

Mark D.M. Leiserson, Matthew A. Reyna, Benjamin J. Raphael

Department of Computer Science and Center for Computational Molecular Biology, Brown University, Providence, 02912, USA

S1 Introduction

In Section S2 of the supplement, we provide a partial derivation of the saddlepoint approximation for the WR-exclusivity p-value. In Section S3 of the supplement, we include additional figures and tables from the experiments in the manuscript, including

- the distributions of sample mutation frequencies for different cancer types,
- a comparison of the R-exclusivity test and the WR-exclusivity test with a weights given by the R-exclusivity test sample space Ω_{R}, and
- the top triples found by the R-exclusivity test and the WR-exclusivity test on thyroid cancer (THCA) with weights given by the RC-exclusivity test sample space Ω_{RC}.

S2 Derivation of the saddlepoint approximation

In this section, we derive a saddlepoint approximation for the WR-exclusivity p-value for a set M of k events. Our approach applies to arbitrary sets M of arbitrary size $k \geq 2$, but this derivation considers $k = 3$ for the ease of presentation.

Our approach is inspired by [3], who derive a saddlepoint approximation for a weighted enrichment test for differentially expressed genes in Gene Ontology categories, and it follows the approach of [1] in greater detail.

As described in the text of the manuscript, for our model, we assume that $\{X_{ij}\}_{j=1}^{n}$ is a set of mutually independent Bernoulli trials for each gene g_i with success probabilities $W = [w_{ij}]$, i.e.,

$$\Pr(X_{ij} = \ell) = \begin{cases} w_{ij}, & \text{if } \ell = 1, \\ 1 - w_{ij}, & \text{if } \ell = 0, \end{cases} \tag{1}$$

where w_{ij} is the probability that gene g_i is mutated in sample s_j. Let $T_{M,j}$ be a random variable with $T_{M,j} = 1$ if s_j has a mutually exclusive mutation in a gene set M and $T_{M,j} = 0$ otherwise. Therefore, $Y_i = \sum_{j=1}^{n} X_{ij}$ is a Poisson binomial distributed variable for the number of mutations in g_i and $T_M = \sum_{j=1}^{n} T_{M,j}$ is a test statistic for mutual exclusivity indicating the number of mutually exclusive mutations in M. We want to find the tail probability (commonly referred to as the p-value) of observing at least t_M mutually exclusive mutations in M given that g_i is mutated in r_i samples. The WR-exclusivity p-value $\Phi_{\mathrm{WR}}(M)$ is the probability of observing at least t_M mutations in a gene set M under this model with

$$\Phi_{\mathrm{WR}}(M) = \Pr(T_M \geq t_M \mid Y_M = \mathbf{r}_M) \tag{2}$$

where $Y_M = [Y_i]_{i \in M}$ and $\mathbf{r}_M = [r_i]_{i \in M}$.

The saddlepoint approximation is given by

$$\Pr(T_M \geq t_M \mid Y_M = \mathbf{r}_M) \approx 1 - \Phi(\tilde{w}) - \phi(\tilde{w}) \left(\frac{1}{\tilde{w}} - \frac{1}{\tilde{u}} \right), \tag{3}$$

where Φ and ϕ are, in this setting, the cumulative distribution and density functions, respectively, of the standard normal distribution, and $\tilde{w}$ and $\tilde{u}$ are defined as follows.

Without loss of generality, let $M = \{1, \ldots, k\}$. First, for $\lambda \in \mathbb{R}^{k+1}$, let $\mathcal{M}_{Y_M, T_M}(\lambda) = E[e^{\sum_{i \in M} \lambda_i Y_i + \lambda_{k+1} T_M}]$ be the joint moment generating function of $\{Y_i\}_{i \in M}$ and T_M, and let $\mathcal{K}_{Y_M, T_M}(\lambda) = \log \mathcal{M}_{Y_M, T_M}(\lambda)$ be the corresponding joint cumulant generating function.

By the mutual independence of $\{Y_i\}_{i\in M}$, we have

$$\begin{aligned}
&K_{Y_M,T_M}(\lambda)\\
&\quad= \log E[e^{\lambda_1 Y_1+\lambda_2 Y_2+\lambda_3 Y_3+\lambda_4 T_M}]\\
&\quad= \log E[e^{\lambda_1 \sum_j X_{1,j}+\lambda_2 \sum_j X_{2,j}+\lambda_3 \sum_j X_{3,j}+\lambda_4 \sum_j T_{M,j}}]\\
&\quad= \log E[e^{\sum_j \lambda_1 X_{1,j}+\lambda_2 X_{2,j}+\lambda_3 X_{3,j}+\lambda_4 T_{M,j}}]\\
&\quad= \log \prod_{j=1}^{n} E[e^{\lambda_1 X_{1,j}+\lambda_2 X_{2,j}+\lambda_3 X_{3,j}+\lambda_4 T_{M,j}}]\\
&\quad= \sum_{j=1}^{n} \log E[e^{\lambda_1 X_{1,j}+\lambda_2 X_{2,j}+\lambda_3 X_{3,j}+\lambda_4 T_{M,j}}].
\end{aligned} \tag{4}$$

Note that $T_{M,j}$ can be written in terms of $\{X_{ij}\}_{i\in M}$. In this case, we have

$$\begin{aligned}
T_{M,j} = {} & X_{1j}(1-X_{2j})(1-X_{3j})\\
& + (1-X_{1j})X_{2j}(1-X_{3j}) + (1-X_{1j})(1-X_{2j})X_{3j}.
\end{aligned} \tag{5}$$

Since $\{X_{ij}\}$ are mutually independent and since the expected value of a Bernoulli random variable is equal to its success probability, we have

$$K_{Y_M,T_M}(\lambda) = \sum_{j=1}^{n} \log L^{(j)}_{Y_M,T_M}(\lambda), \tag{6}$$

where

$$\begin{aligned}
&L^{(j)}_{Y_M,T_M}(\lambda)\\
&\quad= E[e^{\lambda_1 X_{1,j}+\lambda_2 X_{2,j}+\lambda_3 X_{3,j}+\lambda_4 T_{M,j}}]\\
&\quad= (1-w_{1,j})(1-w_{2,j})(1-w_{3,j}) + e^{\lambda_1+\lambda_4} w_{1,j}(1-w_{2,j})(1-w_{3,j})\\
&\qquad + e^{\lambda_2+\lambda_4}(1-w_{1,j})w_{2,j}(1-w_{3,j}) + e^{\lambda_3+\lambda_4}(1-w_{1,j})(1-w_{2,j})w_{3,j}\\
&\qquad + e^{\lambda_1+\lambda_2} w_{1,j}w_{2,j}(1-w_{3,j}) + e^{\lambda_1+\lambda_3} w_{1,j}(1-w_{2,j})w_{3,j}\\
&\qquad + e^{\lambda_2+\lambda_3}(1-w_{1,j})w_{2,j}w_{3,j} + e^{\lambda_1+\lambda_2+\lambda_3} w_{1,j}w_{2,j}w_{3,j}.
\end{aligned} \tag{7}$$

Each of the above terms corresponds to one of the 2^k binary assignments for $\{X_{ij}\}_{i\in M}$, e.g., the term $(1-w_{1,j})(1-w_{2,j})(1-w_{3,j})$ corresponds to assignments $X_{1j} = X_{2j} = X_{3j} = 0$. For the simplicity of notation, let $L^{(j)}(\lambda) = L^{(j)}_{Y_M,T_M}(\lambda)$ for the rest of the derivation.

Similarly, let $\mathcal{M}_{Y_i}(\lambda) = E[e^{\lambda_i Y_i}]$ be the moment generating function of Y_i, and let $\mathcal{K}_{Y_i}(\lambda) = \log \mathcal{M}_{Y_i}(\lambda)$ be the corresponding cumulant generating function. The derivations for these quantities are similar.

Next, let $\mathcal{K}'_{Y_M,T_M}(\lambda)$ be the gradient vector of $\mathcal{K}_{Y_M,T_M}(\lambda)$. For example, we see that

$$\frac{\partial}{\partial\lambda_1} K_{Y_M,T_M}(\lambda) = \sum_{j=1}^{n} \frac{L^{(j)}_{\lambda_1}(\lambda)}{L^{(j)}(\lambda)}, \tag{8}$$

where

$$\begin{aligned} L^{(j)}_{\lambda_1} &= \frac{\partial}{\partial\lambda_1} L^{(j)}_{Y_M,T_M}(\lambda) \\ &= e^{\lambda_1+\lambda_4} w_{1,j}(1-w_{2,j})(1-w_{3,j}) + e^{\lambda_1+\lambda_2} w_{1,j}w_{2,j}(1-w_{3,j}) \\ &\quad + e^{\lambda_1+\lambda_3} w_{1,j}(1-w_{2,j})w_{3,j} + e^{\lambda_1+\lambda_2+\lambda_3} w_{1,j}w_{2,j}w_{3,j}. \end{aligned} \tag{9}$$

The derivations for the remaining entries of $\mathcal{K}'_{Y_M,T_M}(\lambda)$ are similar.

Furthermore, let $\mathcal{K}''_{Y_M,T_M}(\lambda)$ be the Hessian matrix of $\mathcal{K}_{Y_M,T_M}(\lambda)$. For example, we see that

$$\frac{\partial^2}{\partial\lambda_1\partial\lambda_2} K_{Y_M,T_M}(\lambda) = \sum_{j=1}^{n} \frac{L^{(j)}_{\lambda_1\lambda_2}}{L^{(j)}} - \frac{L^{(j)}_{\lambda_1} L^{(j)}_{\lambda_2}}{\left(L^{(j)}\right)^2}, \tag{10}$$

where $L^{(j)}_{\lambda_1}$ and $L^{(j)}_{\lambda_2}$ were already derived for the derivation of $\mathcal{K}'_{Y_M,T_M}$ and

$$\begin{aligned} L^{(j)}_{\lambda_1\lambda_2} &= \frac{\partial^2}{\partial\lambda_1\partial\lambda_2} L^{(j)}_{Y}(\lambda) \\ &= e^{\lambda_1+\lambda_2} w_{1,j}w_{2,j}(1-w_{3,j}) + e^{\lambda_1+\lambda_2+\lambda_3} w_{1,j}w_{2,j}w_{3,j}. \end{aligned} \tag{11}$$

The derivations for the remaining entries of $\mathcal{K}''_{Y_M,T_M}(\lambda)$ are similar.

Finally, define $\tilde{w}$ by

$$\tilde{w} = \sqrt{2}\,\mathrm{sgn}\,(\tilde{y}_{k+1}) \sqrt{\sum_{i\in M} \mathcal{K}_{Y_i}(\hat{x}_i) - \mathcal{K}_{Y_M,T_M}(\tilde{y}) - \tilde{y}^T(\hat{x}-\tilde{x})} \tag{12}$$

and $\tilde{u}$ by

$$\tilde{u} = 2\sinh\left(\frac{\tilde{y}_{k+1}}{2}\right) \sqrt{\frac{\left|\mathcal{K}''_{Y_M,T_M}(\tilde{y})\right|}{\prod_{i\in M} \mathcal{K}''_{Y_i}(\hat{x}_i)}}, \tag{13}$$

where $\tilde{x} = \left(r_1,\ldots,r_k,t_M-\frac{1}{2}\right)$ and $\tilde{y} = (\tilde{y}_1,\ldots,\tilde{y}_{k+1})$ with $\tilde{y}$ the unique solution for $\mathcal{K}'_{Y_M,T_M}(\tilde{y}) = \tilde{x}$ and (3) undefined if $\tilde{y}_{k+1} = 0$, and $\hat{x} = (\hat{x}_1,\ldots,\hat{x}_k,0)$ with $\hat{x}_i$ the unique solution for $\mathcal{K}'_{Y_i}(\hat{x}_i) = r_i$.

Note that, in general, $\mathcal{K}'_{Y_M,T_M}(\tilde{y}) = \tilde{x}$ is a system of $k+1$ nonlinear equations, and $\mathcal{K}'_{Y_i}(\hat{x}_i)$ is a sequence of k nonlinear equations. A root finder, e.g., the MINPACK routine HYBRJ [4], can solve these systems of equations when given expressions for the gradient vector $\mathcal{K}'_{Y_M,T_M}(\lambda)$ and the Hessian matrix $\mathcal{K}''_{Y_M,T_M}(\lambda)$; the MINPACK routine HYBRD approximates $\mathcal{K}''_{Y_M,T_M}(\lambda)$. Various wrappers for HYBRJ and HYBRD exist in various languages, including SciPy's Python wrapper [2], which we used in our code.

S3 Supplemental figures and tables

Figure S1 compares the number of mutated genes per sample in thyroid cancers (THCA), colorectal cancers (COADREAD), and endometrial carcinomas (UCEC), demonstrating high mutation rates in several COADREAD and UCEC samples and high mutational variability in COADREAD and UCEC data.

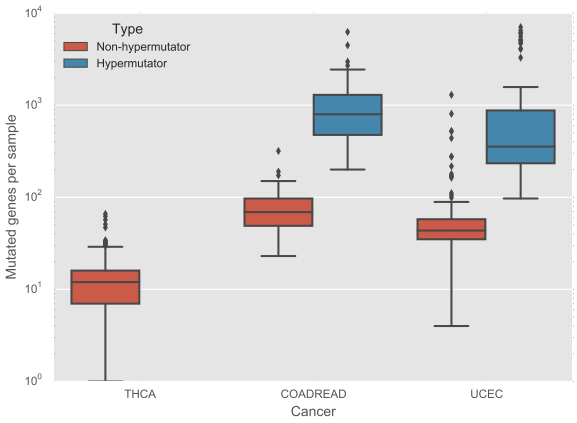

Figure S1: Boxplot of the number of mutated genes per sample in the THCA, COADREAD, and UCEC datasets. The COADREAD and UCEC samples are further broken down into hypermutators (blue) and non-hypermutators (red).

Figure S2 shows a comparison of the p-values and runtimes given by the recusive formula and the saddlepoint approximation methods. On these datasets, the saddlepoint approximation is an extremely accurate approximation of the tail enumeration procedure ($\rho^2 = 0.995$). Additionally, while the median runtimes of the two algorithms are similar, the tail enumeration

procedure is much slower for sets with co-occurring mutations while the saddlepoint approximation is largely unaffected. We expect the discrepancy between runtimes to grow for gene sets of larger sizes. Also, the saddlepoint approximation can be optimized when approximating the R-exclusivity test because each row of the corresponding weight matrix has identical entries, but we did not do so here.

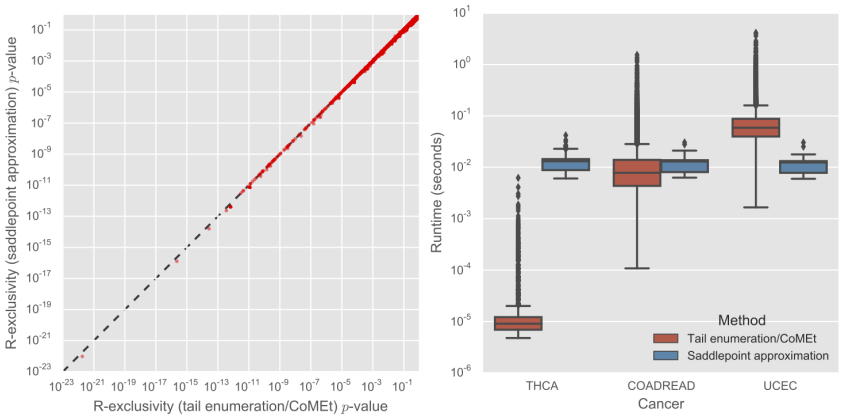

Figure S2: Comparison of the tail enumeration and saddlepoint approximation algorithms for computing the R-exclusivity p-value $\Phi_{\mathrm{R}}(M)$ on triples in THCA, COADREAD, and UCEC. (left) Scatter plot comparing the p-values given by the tail enumeration algorithm (x-axis) versus the saddlepoint approximation (y-axis). (right) Distribution of the runtime (in seconds) required to compute each method on a single triple.

Table S1 shows the five most significant triples identified by the R-exclusivity and WR-exclusivity tests on thyroid cancer (THCA).

$\Phi_{\mathbf{R}}$ **rank**	$\Phi_{\mathbf{WR}}$ **rank**	**Triple** M	$\Phi_{\mathbf{R}}(M)$	$\Phi_{\mathbf{WR}}(M)$
1	1	BRAF, HRAS, NRAS	$1.79 \cdot 10^{-22}$	$1.79 \cdot 10^{-27}$
2	2	BRAF, EIF1AX, NRAS	$2.27 \cdot 10^{-16}$	$2.73 \cdot 10^{-20}$
3	3	**BDP1**, BRAF, NRAS	$2.70 \cdot 10^{-14}$	$6.01 \cdot 10^{-18}$
4	4	BRAF, NRAS, TG	$3.51 \cdot 10^{-13}$	$5.62 \cdot 10^{-17}$
5	5	BRAF, **MUC5B**, NRAS	$6.58 \cdot 10^{-13}$	$1.90 \cdot 10^{-16}$

Table S1: Five most significant triples identified by the R-exclusivity and WR-exclusivity tests on the THCA dataset. Genes in bold are among the 600 longest genes (at least 9,560 amino acids). Seven triples tie for the fifth smallest R-exclusivity p-value, so we indicate the triple with the smallest WR-exclusivity p-value in the table.

References

[1] Ronald W Butler. *Saddlepoint approximations with applications*, volume 22. Cambridge University Press, 2007.

[2] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001–.

[3] David Manescu and Uri Keich. A Symmetric Length-Aware Enrichment Test. *RECOMB*, pages 224–242, 2015.

[4] Jorge J Moré, Burton S Garbow, and Kenneth E Hillstrom. User guide for minpack-1. Technical report, 1980.